arXiv:1512.02671v1 [math.NA] 8 Dec 2015

Householder QR Factorization: Adding Randomization for Column Pivoting

FLAME Working Note #78

Per-Gunnar Martinsson* Gregorio Quintana-Ortí† Nathan Heavner*

Robert van de Geijn‡

November 24, 2015

Abstract

A fundamental problem when adding column pivoting to the Householder QR factorization is that only about half of the computation can be cast in terms of high performing matrix-matrix multiplication, which greatly limits the benefits of so-called blocked algorithms. This paper describes a technique for selecting groups of pivot vectors by means of randomized projections. It is demonstrated that the asymptotic flop count for the proposed method is $2mn^2 - (2/3)n^3$ for an $m \times n$ matrix with $m \geq n$, identical to that of the best classical Householder QR factorization (with or without pivoting). Experiments demonstrate improvements in speed of substantial integer factors (between a factor of 3 and 5) relative to LAPACK's GEQP3 implementation on a modern CPU.

1 Introduction

The QR factorization is a staple of linear algebra, with applications ranging from Linear Least-Squares solution of overdetermined systems to the identification of low rank approximation via the determination of an approximate orthonormal basis for the column space. Standard algorithms for computing the QR factorization include Gram-Schmidt orthogonalization and those based on Householder transformations (reflectors). When it is desirable for the QR factorization to also reveal the approximate rank of the original matrix, it is important that the elements of the diagonal of R be ordered with larger elements in magnitude appearing earlier. In this case, column pivoting (swapping) is employed during the QR factorization, yielding QR factorization with column pivoting (QRP). It is well-known that the Householder QR factorization (HQR) yields columns of Q that are orthogonal to a high degree of precision, making these algorithms the weapon of choice in many situations. Pivoting can be added to HQR to yield HQR with column pivoting (HQRP). This topic is covered by standard texts on numerical linear algebra [11].

To achieve high performance for dense linear algebra algorithms, so-called blocked algorithms are employed that cast most computation in terms of matrix-matrix operations supported by the widely used level-3 Basic Linear Algebra Subprograms (BLAS) [6, 7] because such operations can be implemented to achieve very high performance on modern processors via a combination of careful reuse of data in the caches and low level implementation in terms of assembly code or intrinsic vector operations. Widely used current implementations of the level-3 BLAS are based on techniques exposed by Goto [13, 12] and available in open source libraries including the OpenBLAS [35] (a fork of the GotoBLAS) and BLIS [32], as well as vendor implementations including AMD's ACML [1], Intel's MKL [18], and IBM's ESSL [17] libraries.

*Department of Applied Mathematics, University of Colorado at Boulder, 526 UCB, Boulder, CO 80309-0526, USA

†Depto. de Ingeniería y Ciencia de Computadores, Universitat Jaume I, 12.071 Castellón, Spain

‡Department of Computer Science and Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX 78712 USA.

The fundamental problem with the classical approach to HQRP is that only half of the computation can be cast in terms of the matrix-matrix multiplication (GEMM), as described in the paper [27] that underlies LAPACK's `geqp3` routine [2]. This means that blocking can only improve performance by approximately a factor two, which is inherent from the fact that it must be known how remaining columns will be updated in order to compute the 2-norms of remaining columns. An attempt to overcome this is a pair of papers by Bischof and Quintana-Ortí [5, 4] that describe an HQR algorithm with window pivoting. While much faster than `geqp3`, this approach is more complicated than the method proposed in this paper and never made it into LAPACK.

The present paper proposes to solve the problem of finding blocks of pivot columns by means of randomized projections. To describe the idea, suppose that we seek to determine a set of b good pivot columns in an $m \times n$ matrix A. We then draw a Gaussian random matrix G of size $b \times m$ and form a *sampling matrix* $Y = GA$. Then determine the b pivot columns by executing QRP of the sampling matrix Y. Since Y is small compared to A, this task can be executed very efficiently. (To be precise, one needs to perform a very slight amount of over-sampling, so that Y consists of, say, $b+5$ rows, see Section 3.1.) With this observation, it becomes easy to block the Householder QR factorization with column pivoting. At each iteration of the blocked algorithm, we use the randomized sampling approach to identify a set of b columns that are then moved to the front of the actual matrix, at which point a regular blocked step of HQR can be used to move the computation forward, optionally with additional column pivoting only within a narrow panel of the matrix. Importantly, the sampling matrix can be cheaply downdated rather than recomputed at each step, allowing the performance of the proposed algorithm to asymptotically approach that of a standard blocked HQR implementation without pivoting. Numerical experiments demonstrate the effectiveness of the proposed approach (Section 4.1) and the quality of the resulting QR factorization (Section 4.2).

The idea of using randomized projections to select groups of pivot columns was first published by one of the authors (Martinsson) on ArXiv in May 2015 [22]. A similar insight was independently discovered and published in September 2015 [8] by Jed Duersch and Ming Gu of UC-Berkeley, also on ArXiv. The technique for downdating of the sampling matrix was discovered independently by the two groups, but was first published in [8]. The method based on random projections that we use here to select pivot columns has been used previously [10, 21, 16, 23, 24] for solving the problem of computing an approximate low-rank approximation to a given matrix. In particular, [23, 21] demonstrate that the randomized method is an efficient and accurate tool for finding a set of columns of a matrix that form a good basis for the column space, which is a problem that is essentially equivalent to finding good pivot columns. This problem is also immediately connected to the problem of finding a set of columns of maximal spanning volume [14], and to the problem of finding so called *CUR* and *interpolative* decompositions [33].

We make a number of new contributions. We propose an algorithm that uses randomization to achieve unprecedented performance for computing HQRP. The algorithm is presented in a way that reflects how high-performance implementations of HQR are implemented in practice in libraries like LAPACK and `libflame` [30, 31] so that the results can be independently reproduced. A cost analysis of the algorithm shows that asymptotically the number of floating point operations approaches that of HQR without pivoting while most computation is cast in terms of matrix-matrix multiplication like the corresponding blocked HQR algorithm. An implementation shows that in practice excellent performance improvements are achieved relative to `geqp3`. This implementation will become part of the `libflame` library. The quality of the computed QR factorization is empirically compared to that of the SVD and `geqp3`. Finally, in the conclusion we discuss future opportunities.

2 Householder QR Factorization

In this section, we briefly review the state-of-the art regarding Householder factorization based on Householder transformations (HQR). Throughout, we use the FLAME notation for representing dense linear algebra algorithms that was first introduced almost two decades ago [26, 15].

2.1 Householder transformations (reflectors)

A standard topic in numerical linear algebra is the concept of a reflector, also known as a Householder transformation [11]. The review in this subsection follows [19] in which a similar notation is also employed.

Given a nonzero vector $u \in \mathbb{C}^n$, the matrix $H(u) = I - \frac{1}{\tau} uu^H$ with $\tau = \frac{u^H u}{2}$ has the property that it reflects a vector to which it is applied with respect to the subspace orthogonal to u. Given a vector x, the vector u and scalar τ can be chosen so that $H(u)x$ equals a multiple of e_0, the first column of the identiy matrix. Furthermore, u can be normalized so that its first element equals one.

In our discussions, given a vector $x = \begin{pmatrix} \chi_1 \\ x_2 \end{pmatrix}$, the function $\left[\begin{pmatrix} \rho \\ u_2 \end{pmatrix}, \tau \right] := \text{Housev} \left(\begin{pmatrix} \chi_1 \\ x_2 \end{pmatrix} \right)$ computes the vector $u = \begin{pmatrix} 1 \\ u_2 \end{pmatrix}$ and $\tau = \frac{u^H u}{2}$ so that $H(u)x = \rho e_0$,

2.2 Unblocked Householder QR factorization

A standard unblocked algorithm for HQR of a given matrix $A \in \mathbb{C}^{m \times n}$, typeset using the FLAME notation, is given in Fig. 1 (left). The body of the loop computes

$$\left[\begin{pmatrix} \rho_{11} \\ u_{21} \end{pmatrix}, \tau_{11} \right] := \text{Housev} \left(\begin{pmatrix} \alpha_{11} \\ a_{21} \end{pmatrix} \right),$$

which overwrites a_{11} with ρ_{11} and a_{21} with u_{21}, after which the remainder of A is updated by

$$\begin{pmatrix} a_{12}^T \\ A_{22} \end{pmatrix} := \left(I - \frac{1}{\tau_{11}} \begin{pmatrix} 1 \\ u_{21} \end{pmatrix} \begin{pmatrix} 1 \\ u_{21} \end{pmatrix}^H \right) \begin{pmatrix} a_{12}^T \\ A_{22} \end{pmatrix}.$$

Upon completion, the (Householder) vectors that define the Householder transformations have overwritten the elements in which they introduced zeroes, and the upper triangular part of A contains R. How the matrix T fits into the picture will become clear next.

2.3 The UT transform: Accumulating Householder transformations

Given $A \in \mathbb{C}^{n \times b}$, let U contain the Householder vectors computed during the HQR of A. Let us assume that $H(u_{b-1}) \cdots H(u_1) H(u_0) A = R$. Then there exists an upper triangular matrix so that $I - UT^{-H}U^H = H(u_{b-1}) \cdots H(u_1) H(u_0)$. The desired matrix T equals the strictly upper triangular part of $U^H U$ with the diagonal elements equal to $\tau_0, \ldots, \tau_{b-1}$. The matrix T can be computed during the unblocked HQR, as indicated in Fig. 1 (left). In [19] we call the transformation $I - UT^{-1}U^H$ that equals the accumulated Householder transformations the *UT transform*. In that paper, we point out similar published results [34, 25], including how this result relates to the more familiar WY transform [3] and compact WY transform [28].

2.4 A blocked QR Householder factorization algorithm

A blocked algorithm for HQR that exploits the insights that resulted in the UT transform can now be described as follows. Partition

$$A \rightarrow \left(\begin{array}{c|c} A_{11} & A_{12} \\ \hline A_{21} & A_{22} \end{array} \right)$$

where A_{11} is $b \times b$. We can use the unblocked algorithm in Fig. 1 (left) to factor the panel $\begin{pmatrix} A_{11} \\ A_{21} \end{pmatrix}$, creating matrix T_{11} as a side effect. Now we need to also apply the UT transform to the rest of the columns:

$$\begin{pmatrix} A_{12} \\ A_{22} \end{pmatrix} := \left(I - \begin{pmatrix} U_{11} \\ U_{21} \end{pmatrix} T_{11}^{-1} \begin{pmatrix} U_{11} \\ U_{21} \end{pmatrix}^H \right)^H \begin{pmatrix} A_{12} \\ A_{22} \end{pmatrix} = \begin{pmatrix} A_{12} - U_{11} W_{12} \\ A_{22} - U_{21} W_{12} \end{pmatrix},$$

Algorithm: $[A, T] := \text{HQR_UNB}(A, T)$

Partition $A \rightarrow \left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right)$,
$T \rightarrow \left(\begin{array}{c|c} T_{TL} & T_{TR} \\ \hline 0 & T_{BR} \end{array}\right)$
where A_{TL} is 0×0, T_{TL} is 0×0

while $m(A_{TL}) < m(A)$ **do**

Repartition

$$\left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right) \rightarrow \left(\begin{array}{c|c|c} A_{00} & a_{01} & A_{02} \\ \hline a_{10}^T & \alpha_{11} & a_{12}^T \\ \hline A_{20} & a_{21} & A_{22} \end{array}\right),$$
$$\left(\begin{array}{c|c} T_{TL} & T_{TR} \\ \hline 0 & T_{BR} \end{array}\right) \rightarrow \left(\begin{array}{c|c|c} T_{00} & t_{01} & T_{02} \\ \hline 0 & \tau_{11} & t_{12}^T \\ \hline 0 & 0 & T_{22} \end{array}\right)$$

where α_{11} is 1×1, τ_{11} is 1×1

$$\left[\begin{pmatrix} \alpha_{11} \\ a_{21} \end{pmatrix}, \tau_{11}\right] := \text{Housev}\begin{pmatrix} \alpha_{11} \\ a_{21} \end{pmatrix}$$
$$w_{12}^T := (a_{12}^T + a_{21}^H A_{22})/\tau_{11}$$
$$\begin{pmatrix} a_{12}^T \\ A_{22} \end{pmatrix} := \begin{pmatrix} a_{12}^T - w_{12}^T \\ A_{22} - a_{21} w_{12}^T \end{pmatrix}$$
$$t_{01} := (a_{10}^T)^H + A_{20}^H a_{21}$$

Continue with

$$\left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right) \leftarrow \left(\begin{array}{c|c|c} A_{00} & a_{01} & A_{02} \\ \hline a_{10}^T & \alpha_{11} & a_{12}^T \\ \hline A_{20} & a_{21} & A_{22} \end{array}\right),$$
$$\left(\begin{array}{c|c} T_{TL} & T_{TR} \\ \hline 0 & T_{BR} \end{array}\right) \leftarrow \left(\begin{array}{c|c|c} T_{00} & t_{01} & T_{02} \\ \hline 0 & \tau_{11} & t_{12}^T \\ \hline 0 & 0 & T_{22} \end{array}\right)$$

endwhile

Algorithm: $[A, T] := \text{HQR_BLK}(A, T)$

Partition $A \rightarrow \left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right)$,
$T \rightarrow \left(\begin{array}{c} T_T \\ \hline T_B \end{array}\right)$
where A_{TL} is 0×0, T_T has 0 rows

while $m(A_{TL}) < m(A)$ **do**

Determine block size b

Repartition

$$\left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right) \rightarrow \left(\begin{array}{c|c|c} A_{00} & A_{01} & A_{02} \\ \hline A_{10} & A_{11} & A_{12} \\ \hline A_{20} & A_{21} & A_{22} \end{array}\right),$$
$$\left(\begin{array}{c} T_T \\ \hline T_B \end{array}\right) \rightarrow \left(\begin{array}{c} T_0 \\ \hline T_1 \\ \hline T_2 \end{array}\right)$$

where A_{11} is $b \times b$, T_1 has b rows

$$\left[\begin{pmatrix} A_{11} \\ A_{21} \end{pmatrix}, T_1\right] :=$$
$$\text{HQR_UNB}\left(\begin{pmatrix} A_{11} \\ A_{21} \end{pmatrix}, T_1\right)$$
$$W_{12} := T_1^{-H}(U_{11}^H A_{12} + U_{21}^H A_{22})$$
$$\begin{pmatrix} A_{12} \\ A_{22} \end{pmatrix} := \begin{pmatrix} A_{12} - U_{11} W_{12} \\ A_{22} - U_{21} W_{12} \end{pmatrix}$$

Continue with

$$\left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right) \leftarrow \left(\begin{array}{c|c|c} A_{00} & A_{01} & A_{02} \\ \hline A_{10} & A_{11} & A_{12} \\ \hline A_{20} & A_{21} & A_{22} \end{array}\right),$$
$$\left(\begin{array}{c} T_T \\ \hline T_B \end{array}\right) \leftarrow \left(\begin{array}{c} T_0 \\ \hline T_1 \\ \hline T_2 \end{array}\right)$$

endwhile

Figure 1: Left: Unblocked Householder transformation based QR factorization merged with the computation of T for the UT transform. Right: Blocked Householder transformation based QR factorization. In this algorithm, U_{11} is the unit lower triangular matrix stored below the diagonal of A_{11} and U_{21} is stored in A_{21}.

where $W_{12} = T_{11}^{-H}(U_{11}^H A_{12} + U_{21}^H A_{22})$. This motivates the blocked HQR algorithm in Fig. 1 (right) which we will refer to as HQR_BLK.

The benefit of the blocked algorithm is that it casts most computation in terms of the computations $U_{21}^H A_{22}$ (row panel times matrix multiply) and $A_{22} - U_{21} W_{12}$ (rank-b update). Such matrix-matrix multiplications can attain high performance by amortizing data movement between memory layers.

These insights form the basis for the LAPACK routine GEQRF (except that it uses a compact WY transform instead of the UT transform).

Algorithm: $[A, t, p] = \text{HQRP_UNB_VAR1}\ (A)$

v := ComputeWeights(A)

Partition $A \rightarrow \left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right)$, $t \rightarrow \left(\begin{array}{c} t_T \\ \hline t_B \end{array}\right)$, $p \rightarrow \left(\begin{array}{c} p_T \\ \hline p_B \end{array}\right)$, $v \rightarrow \left(\begin{array}{c} v_T \\ \hline v_B \end{array}\right)$

where A_{TL} is 0×0 and t_T, p_T, v_T have 0 elements

while $n(A_{TL}) < n(A)$ **do**

Repartition

$$\left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right) \rightarrow \left(\begin{array}{c|c|c} A_{00} & a_{01} & A_{02} \\ \hline a_{10}^T & \alpha_{11} & a_{12}^T \\ \hline A_{20} & a_{21} & A_{22} \end{array}\right), \left(\begin{array}{c} t_T \\ \hline t_B \end{array}\right) \rightarrow \left(\begin{array}{c} t_0 \\ \hline \tau_1 \\ \hline t_2 \end{array}\right), \left(\begin{array}{c} p_T \\ \hline p_B \end{array}\right) \rightarrow \left(\begin{array}{c} p_0 \\ \hline \pi_1 \\ \hline p_2 \end{array}\right), \left(\begin{array}{c} v_T \\ \hline v_B \end{array}\right) \rightarrow \left(\begin{array}{c} v_0 \\ \hline \nu_1 \\ \hline v_2 \end{array}\right)$$

where α_{11}, τ_1, π_1, and ν_1 are scalars

$$[\left(\begin{array}{c} \nu_1 \\ \hline v_2 \end{array}\right), \pi_1] = \text{DeterminePivot}(\left(\begin{array}{c} \nu_1 \\ \hline v_2 \end{array}\right))$$

$$\left(\begin{array}{c|c} a_{01} & A_{02} \\ \hline \alpha_{11} & a_{12}^T \\ \hline a_{21} & A_{22} \end{array}\right) := \left(\begin{array}{c|c} a_{01} & A_{02} \\ \hline \alpha_{11} & a_{12}^T \\ \hline a_{21} & A_{22} \end{array}\right) P(\pi_1)^T$$

$$\left[\left(\begin{array}{c} \alpha_{11} \\ \hline a_{21} \end{array}\right), \tau_1\right] := \left[\left(\begin{array}{c} \rho_{11} \\ \hline u_{21} \end{array}\right), \tau_1\right] = \text{Housev}\left(\begin{array}{c} \alpha_{11} \\ \hline a_{21} \end{array}\right)$$

$$w_{12}^T := (a_{12}^T + a_{21}^H A_{22})/\tau_1$$

$$\left(\begin{array}{c} a_{12}^T \\ \hline A_{22} \end{array}\right) := \left(\begin{array}{c} a_{12}^T - w_{12}^T \\ \hline A_{22} - a_{21} w_{12}^T \end{array}\right)$$

$$v_2 = \text{UpdateWeights}(v_2, a_{12})$$

Continue with

...

endwhile

Figure 2: Simple unblocked QRP via Householder transformations (HQRP Unblocked Variant 1).

2.5 Householder QR factorization with column pivoting

A (rank-revealing) QR factorization with column pivoting (QRP) computes $AP = QR$ where P is a permutation matrix that swaps the columns $of A$ so that the diagonal elements of R are (approximately) ordered from largest to smallest (in magnitude). If we partition

$$Q = \left(\begin{array}{c|c} Q_L & Q_R \end{array}\right) \quad \text{and (conformally)} \quad R = \left(\begin{array}{c|c} R_{TL} & R_{TR} \\ \hline 0 & R_{BR} \end{array}\right),$$

where Q_L has k columns. Then the column space of Q_L equals approximately the column space of A and

$$A \approx Q_L \left(\begin{array}{c|c} R_{TL} & R_{TR} \end{array}\right) P^T,$$

which shows how identifying Q_L and part of R can be used to create a low rank approximation for A.

We start by describing the simplest way of adding column pivoting to the algorithm in Fig. 1, ignoring the computation of T in that algorithm, resulting in the algorithm in Fig. 2. Which column to pivot to the front is determined by computing the 2-norms (Euclidean lengths) of the columns of A. For convenience, we

Algorithm:
$[A, T, s, v, W] = \text{HQRP_UNB}\ (A, T, s, v, W, r)$

Partition $A \rightarrow \left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right)$,
$T \rightarrow \left(\begin{array}{c|c} T_{TL} & T_{TR} \\ \hline T_{BL} & T_{BR} \end{array}\right)$, $s \rightarrow \left(\begin{array}{c} s_T \\ \hline s_B \end{array}\right)$,
$v \rightarrow \left(\begin{array}{c} v_T \\ \hline v_B \end{array}\right)$, $W \rightarrow \left(\begin{array}{c|c} W_{TL} & W_{TR} \\ \hline W_{BL} & W_{BR} \end{array}\right)$
where A_{TL}, T_{TL}, W_{TL} are 0×0,
v_T has 0 elements

while $n(A_{TL}) < r$ **do**

Repartition

$$\left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right) \rightarrow \left(\begin{array}{c|c|c} A_{00} & a_{01} & A_{02} \\ \hline a_{10}^T & \alpha_{11} & a_{12}^T \\ \hline A_{20} & a_{21} & A_{22} \end{array}\right),$$

$$\left(\begin{array}{c|c} T_{TL} & T_{TR} \\ \hline T_{BL} & T_{BR} \end{array}\right) \rightarrow \left(\begin{array}{c|c|c} T_{00} & t_{01} & T_{02} \\ \hline 0 & \tau_{11} & t_{12}^T \\ \hline 0 & 0 & T_{22} \end{array}\right),$$

$$\left(\begin{array}{c} s_T \\ \hline s_B \end{array}\right) \rightarrow \left(\begin{array}{c} s_0 \\ \hline \sigma_1 \\ \hline s_2 \end{array}\right), \left(\begin{array}{c} v_T \\ \hline v_B \end{array}\right) \rightarrow \left(\begin{array}{c} v_0 \\ \hline \nu_1 \\ \hline v_2 \end{array}\right),$$

$$\left(\begin{array}{c|c} W_{TL} & W_{TR} \\ \hline W_{BL} & W_{BR} \end{array}\right) \rightarrow \left(\begin{array}{c|c|c} W_{00} & w_{01} & W_{02} \\ \hline w_{10}^T & \omega_{11} & w_{12}^T \\ \hline W_{20} & w_{21} & W_{22} \end{array}\right)$$

where α_{11}, τ_{11}, $\cdots$ are scalars

$$\left[\left(\begin{array}{c} \nu_1 \\ \hline v_2 \end{array}\right), \sigma_1\right] = \text{DeterminePivot}\left(\left(\begin{array}{c} \nu_1 \\ \hline v_2 \end{array}\right)\right)$$

$$\left(\begin{array}{c|c} a_{01} & A_{02} \\ \hline \alpha_{11} & a_{12}^T \\ \hline a_{21} & A_{22} \end{array}\right) := \text{Swap}\left(\sigma_1, \left(\begin{array}{c|c} a_{01} & A_{02} \\ \hline \alpha_{11} & a_{12}^T \\ \hline a_{21} & A_{22} \end{array}\right)\right)$$

$\alpha_{11} := \alpha_{11} - a_{10}^T w_{01}$
$a_{21} := a_{21} - A_{20} w_{01}$
$a_{12}^T := a_{12}^T - a_{10}^T W_{02}$
$\left[\left(\begin{array}{c} \alpha_{11} \\ \hline a_{21} \end{array}\right), \tau_{11}\right] := \text{Housev}\left(\begin{array}{c} \alpha_{11} \\ \hline a_{21} \end{array}\right)$
$w_{12}^T := (a_{12}^T + a_{21}^H A_{22} - (a_{21}^H A_{20}) W_{02})/\tau_{11}$
$a_{12}^T := a_{12}^T - w_{12}^T$
$v_2 = \text{DowndateWeights}(v_2, a_{12})$

Continue with

...

endwhile

Algorithm:
$[A, T, s] := \text{HQRP_BLK}(A, T, s, r)$

$v := \text{ComputeWeights}(A)$

Partition $A \rightarrow \left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right)$,
$T \rightarrow \left(\begin{array}{c} T_T \\ \hline T_B \end{array}\right)$, $s \rightarrow \left(\begin{array}{c} s_T \\ \hline s_B \end{array}\right)$,
$v \rightarrow \left(\begin{array}{c} v_T \\ \hline v_B \end{array}\right)$, $W \rightarrow \left(\begin{array}{c|c} W_L & W_R \end{array}\right)$
where A_{TL} is 0×0; W_L has 0 cols;
T_T, s_T, and v_T have 0 rows

while $n(A_{TL}) < r$) **do**
Determine block size b
Repartition

$$\left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right) \rightarrow \left(\begin{array}{c|c|c} A_{00} & A_{01} & A_{02} \\ \hline A_{10} & A_{11} & A_{12} \\ \hline A_{20} & A_{21} & A_{22} \end{array}\right),$$

$$\left(\begin{array}{c} T_T \\ \hline T_B \end{array}\right) \rightarrow \left(\begin{array}{c} T_0 \\ \hline T_1 \\ \hline T_2 \end{array}\right),$$

$$\left(\begin{array}{c} s_T \\ \hline s_B \end{array}\right) \rightarrow \left(\begin{array}{c} s_0 \\ \hline s_1 \\ \hline s_2 \end{array}\right), \left(\begin{array}{c} v_T \\ \hline v_B \end{array}\right) \rightarrow \left(\begin{array}{c} v_0 \\ \hline v_1 \\ \hline v_2 \end{array}\right),$$

$$\left(\begin{array}{c|c} W_L & W_R \end{array}\right) \rightarrow \left(\begin{array}{c|c|c} W_0 & W_1 & W_2 \end{array}\right)$$

where A_{11} is $b \times b$, W_1 has b cols,
T_1, s_1, v_1 have b rows

$$\left[\left(\begin{array}{c|c} A_{11} & A_{12} \\ \hline A_{21} & A_{22} \end{array}\right), T_1, s_1, \left(\begin{array}{c} v_1 \\ \hline v_2 \end{array}\right), \left(\begin{array}{c|c} W_1 & W_2 \end{array}\right)\right]$$
$:= \text{HQRP_UNB}$
$$\left(\left(\begin{array}{c|c} A_{11} & A_{12} \\ \hline A_{21} & A_{22} \end{array}\right), T_1, s_1, \left(\begin{array}{c} v_1 \\ \hline v_2 \end{array}\right), \left(\begin{array}{c|c} W_1 & W_2 \end{array}\right), b\right)$$

$A_{22} := A_{22} - A_{21} W_2$

Continue with

...

endwhile

Figure 3: Left: Unblocked HQRP. Right: Blocked HQRP.

will use the squares of the 2-norms and compute these as the vector $v = \text{ComputeWeights}(A)$, where the ith entry in v initially equals the weight (square of the 2-norms)* of the ith column of A. In the first step, we then determine the largest entry in v, permute the corresponding column of A to the front (swapping it with the first column). Then, the algorithm proceeds as does HQR_UNB in Fig. 1. At the bottom of the loop, rather then recomputing v for A_{22} (which becomes the matrix A_{BR} with which the algorithm will continue in the next iteration), we observe that the values of v_2 can be updated by subtracting the squares of the corresponding elements of the new row of R that has been computed, r_{12}^T, which has overwritten the entries of a_{12}^T. In other words, rather than recomputing v_2, that vector can be *downdated* with the entries of r_{12}, thus saving an amount of computation that is proportional to the other costs in this iteration. What is important is that in order to determine how columns need to be swapped, the vector v_2 must be updated in each iteration and to update the vector v_2, the elements of r_{12}^T must be computed. That, in turn, means that the vector w_{12}^T must be computed. This limits how much computation can be cast in terms of matrix-matrix multiplication, as we will see next.

The important observation is that at each step the unblocked algorithm must at the very least update a_{12}^T with r_{12}^T (the next row of R) in each iteration. In this process, the next column of the matrix W that appears in the blocked algorithm in Fig. 1 is computed, and hence there is no point in recomputing it in the blocked algorithm. This leads to the blocked algorithm in Fig. 3 (right), which calls the unblocked algorithm to its left. In HQRP_UNB, the observation is that the updating of α_{11}, a_{21}, and a_{12}^T requires the original contents of those parts of A, the original contents of A_{22}, and the already updated contents of a_{01},A_{20}, and w_{01}. Once α_{11}, a_{21}, and a_{12}^T have been updated, w_{12}^T and r_{12}^T (which overwrites a_{12}^T) can be computed, and v_2 can be downdated. Upon completion, A_{BR} has only been partially updated, requiring the rank-b update $A_{22} := A_{22} - A_{21}W_{12}$ in the blocked algorithm, which then casts only half the computation in terms of GEMM.

The insights we just described were first published in [27] and underlie the routine `geqp3` in LAPACK.

3 Randomization to the Rescue

This section describes a technique for picking, in a computationally efficient manner, a selection of b columns from a given $n \times n$ matrix A that form good choices for the first b pivots in a blocked HQRP algorithm. Observe that this task is mathematically equivalent to finding an index set I of length b such that the columns in $A(:, I)$ form an optimal approximate basis for the column space of A. Another way of expressing this problem is that we are looking for the collection of b columns whose spanning volume in $\mathbb{C}^n$ is maximal. To find the absolutely optimal choice here is a hard problem [14], but luckily, for pivoting purposes it is sufficient to find a choice that is "good enough."

3.1 The basic idea

The strategy that we propose is simple. The idea is to perform classical QR factorization with column pivoting (QRP) on a matrix Y that is much smaller than A, so that performing QRP with that matrix constitutes a lower order cost. As a bonus, it may fit in fast cache memory. This smaller matrix can be constructed by forming random linear combinations of the rows of A as follows:

1. Fix an over-sampling parameter p. Setting $p = 5$ or $p = 10$ are good choices.
2. Form a random matrix G of size $(b + p) \times n$ whose entries are drawn independently from a normalized Gaussian distribution.
3. Form the $(b + p) \times n$ sampling matrix $Y = GA$.

The sampling matrix Y has as many columns as A, but much fewer rows. Now execute b steps of a column pivoted QR factorization to determine an integer vector with b elements that capture how columns need to

*In practice, computing the square of the 2-norms may cause overflow, which is a detail that can be fixed via standard techniques.

Algorithm: $[A, T, s] := \text{HQRRP_BLK}(A, T, s, n_b, p)$

$G := \text{RAND_IID}(n_b + p, n(A))$
$Y := GA$
Partition $A \rightarrow \left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right)$, $T \rightarrow \left(\begin{array}{c} T_T \\ \hline T_B \end{array}\right)$, $s \rightarrow \left(\begin{array}{c} s_T \\ \hline s_B \end{array}\right)$, $Y \rightarrow \left(\begin{array}{c|c} Y_L & Y_R \end{array}\right)$
where A_{TL} is 0×0, T_T has 0 rows, s_T has 0 rows, Y_L has 0 columns

while $m(A_{TL}) < m(A)$ **do**

Determine block size $b = \min(n_b, n(A_{BR}))$
Repartition

$$\left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right) \rightarrow \left(\begin{array}{c|c|c} A_{00} & A_{01} & A_{02} \\ \hline A_{10} & A_{11} & A_{12} \\ \hline A_{20} & A_{21} & A_{22} \end{array}\right), \left(\begin{array}{c} T_T \\ \hline T_B \end{array}\right) \rightarrow \left(\begin{array}{c} T_0 \\ \hline T_1 \\ \hline T_2 \end{array}\right), \left(\begin{array}{c} s_T \\ \hline s_B \end{array}\right) \rightarrow \left(\begin{array}{c} s_0 \\ \hline s_1 \\ \hline s_2 \end{array}\right),$$
$$\left(\begin{array}{c|c} Y_L & Y_R \end{array}\right) \rightarrow \left(\begin{array}{c|c|c} Y_0 & Y_1 & Y_2 \end{array}\right)$$
where A_{11} is $b \times b$, T_1 has b rows, s_1 has b rows, Y_1 has b columns

---

$s_1 := \text{DETERMINEPIVOTS}(\left(\begin{array}{c|c} Y_1 & Y_2 \end{array}\right), b)$
$\left(\begin{array}{c|c} A_{01} & A_{02} \\ \hline A_{11} & A_{12} \\ \hline A_{21} & A_{22} \end{array}\right) := \text{SWAPCOLS}(s_1, \left(\begin{array}{c|c} A_{01} & A_{02} \\ \hline A_{11} & A_{12} \\ \hline A_{21} & A_{22} \end{array}\right))$
$\left[\left(\begin{array}{c} \blacksquare \\ \hline A_{21} \end{array}\right), T_1\right] := \text{HQR_UNB}(\left(\begin{array}{c} A_{11} \\ \hline A_{21} \end{array}\right), T_1)$
$W_{12} := T_1^{-H}(U_{11}^H A_{12} + U_{21}^H A_{22})$
$\left(\begin{array}{c} \blacksquare \\ \hline A_{22} \end{array}\right) := \left(\begin{array}{c} A_{12} - U_{11} W_{12} \\ \hline A_{22} - U_{21} W_{12} \end{array}\right)$
$(\blacksquare \; Y_2) := \text{SWAPCOLS}(s_1, \left(\begin{array}{c|c} Y_1 & Y_2 \end{array}\right))$
$Y_2 := Y_2 - \left(G_1 - (G_1 U_{11} + G_2 U_{21}) T_{11}^{-1} {U_{11}}^H\right) R_{12}.$

Continue with

$$\left(\begin{array}{c|c} A_{TL} & A_{TR} \\ \hline A_{BL} & A_{BR} \end{array}\right) \leftarrow \left(\begin{array}{c|c|c} A_{00} & A_{01} & A_{02} \\ \hline A_{10} & A_{11} & A_{12} \\ \hline A_{20} & A_{21} & A_{22} \end{array}\right), \left(\begin{array}{c} T_T \\ \hline T_B \end{array}\right) \leftarrow \left(\begin{array}{c} T_0 \\ \hline T_1 \\ \hline T_2 \end{array}\right), \left(\begin{array}{c} s_T \\ \hline s_B \end{array}\right) \leftarrow \left(\begin{array}{c} s_0 \\ \hline s_1 \\ \hline s_2 \end{array}\right),$$
$$\left(\begin{array}{c|c} Y_L & Y_R \end{array}\right) \leftarrow \left(\begin{array}{c|c|c} Y_0 & Y_1 & Y_2 \end{array}\right)$$

endwhile

Figure 4: Blocked Householder transformation based QR factorization with column pivoting based on randomization. In this algorithm, U_{11} equals the unit lower triangular matrix stored below the diagonal of A_{11}, $U_{21} = A_{21}$, and $R_{12} = A_{12}$. The steps highlighted in gray constitute the blocked QR factorization without column pivoting from Fig. 1.

be pivoted:

$$s = \text{DETERMINEPIVOTS}(Y, b).$$

In other words, the columns $Y(:, s)$ are good pivot columns for Y. Our claim is that due to the way Y is constructed, the columns $A(:, s)$ are then also good choices for pivot columns of A. This claim is supported by extensive numerical experiments, some of which are given in Section 4.2. There is theory supporting the claim that these b columns form a good approximate basis for the column space of A, see, e.g. [16, Sec. 5.2] and [21, 33], but it has not been directly proven that they form optimal choices as pivots in a QR factorization. This should not be surprising given that it is *known* that even classical column pivoting can result in poor choices [20]. Known algorithms that are provably good are all far more complex to implement [14].

Notice that there are many choices of algorithms that can be employed to determine the pivots. For example, since high numerical accuracy is not necessary, the classical Modified Gram-Schmidt (MGS) algorithm with column pivoting is a simple yet effective choice.

Remark 1 (Choice of over-sampling parameter p) *The reliability of the procedures described in this section depend on the choice of the over-sampling parameter p. It is well understood how large p needs to be in order to determine a high-quality approximate basis for the column space of A with extremely high reliability: the choice $p = 5$ is very good, $p = 10$ is excellent, and $p = b$ is almost always over-kill [16]. The pivot selection problem is less well studied, but is more forgiving. (The choice of pivots does not have to be particularly optimal, it just has to avoid the very bad choices.) Numerical experiments indicate that even setting $p = 0$ typically results in good choices. However, the choices $p = 5$ or $p = 10$ appear to be good generic values that have resulted in excellent choices in every experiment we have run.*

3.2 Efficient downdating of the sampling matrix Y

For the QRP factorization algorithm, it is well known that one does not need to recompute the column norms of the remainder matrix after each step. Instead, these can cheaply be downdated, as described, e.g., in [29, Ch.5, Sec. 2.1]. In terms of asymptotic flop counts, this observation makes the cost of pivoting become a lower order term, and consequently both unpivoted and pivoted Householder QR algorithms have the same leading order term $(4/3)n^3$ in their asymptotic flop counts[†]. In this section, we describe an analogous technique for the randomized sampling strategy described in Section 3.1. This downdating strategy was discovered by one of the authors and also independently by others who first published it to arXiv in September 2015 [8].

First observe that if we use the randomized sampling technique described in Section 3.1 in the most obvious fashion, we would need to draw, at each step of the iteration, a Gaussian random matrix G, and then form a sampling matrix Y via a matrix-matrix multiplication involving the remaining portion of A in the lower right corner. The cost of all these matrix-matrix multiplications would add up to an $O(n^3)$ term. However, it turns out to be possible to avoid recomputing G and Y from scratch at every step. The idea is that if we use the right randomizing matrix G, then the corresponding sampling matrix Y can inexpensively be computed by downdating the sampling matrix from the previous step.

To describe the downdating procedure in detail, let us assume that one iteration of the blocked algorithm has completed, so that, at the bottom of the loop body, the matrix A contains

$$\left(\begin{array}{c|c} A_{11} & A_{12} \\ \hline A_{21} & A_{22} \end{array}\right) = \left(\begin{array}{c|c} U \backslash R_{11} & R_{12} \\ \hline U_{21} & \widehat{A}_{22} - U_{21} W_{12} \end{array}\right).$$

Here $\widehat{A}$ denotes the original contents of AP_1, where P_1 captures how columns have been swapped so far. Hence

$$\left(I - \begin{pmatrix} U_{11} \\ U_{21} \end{pmatrix} T_{11}^{-H} \begin{pmatrix} U_{11} \\ U_{21} \end{pmatrix}^H\right) \left(\begin{array}{c|c} \widehat{A}_{11} & \widehat{A}_{12} \\ \hline \widehat{A}_{21} & \widehat{A}_{22} \end{array}\right) P_1 = \left(\begin{array}{c|c} R_{11} & R_{12} \\ \hline 0 & A_{22} \end{array}\right)$$

[†]Here we count a multiply and add each as one flop, regardless of whether a complex or real operation is performed.

Now, let $\left(\widetilde{G}_1 \middle| \widetilde{G}_2\right)$ denote the next sampling matrix, so that

$$\left(\widetilde{Y}_1 \middle| \widetilde{Y}_2\right) = \left(\widetilde{G}_1 \middle| \widetilde{G}_2\right) \left(\begin{array}{c|c} 0 & 0 \\ \hline 0 & A_{22} \end{array}\right)$$

Notice that the careful placement of zeroes ensures that this is equivalent to choosing $\widetilde{Y}_2 = \widetilde{G}_2 A_{22}$ for some sampling matrix G_2. Letting

$$\left(Q_1 \middle| Q_2\right) = \left(I - \left(\frac{U_{11}}{U_{21}}\right) T_{11}^{-1} \left(\frac{U_{11}}{U_{21}}\right)^H\right)$$

we find that

$$\begin{aligned}
\left(\widetilde{G}_1 \middle| \widetilde{G}_2\right) \left(\begin{array}{c|c} 0 & 0 \\ \hline 0 & A_{22} \end{array}\right) &= \left(\widetilde{G}_1 \middle| \widetilde{G}_2\right) \left(\left(\begin{array}{c|c} R_{11} & R_{12} \\ \hline 0 & A_{22} \end{array}\right) - \left(\begin{array}{c|c} R_{11} & R_{12} \\ \hline 0 & 0 \end{array}\right)\right) \\
&= \left(\widetilde{G}_1 \middle| \widetilde{G}_2\right) \underbrace{\left(Q_1 \middle| Q_2\right)^H \left(Q_1 \middle| Q_2\right)}_{I} \left(\left(\begin{array}{c|c} R_{11} & R_{12} \\ \hline 0 & A_{22} \end{array}\right) - \left(\begin{array}{c|c} R_{11} & R_{12} \\ \hline 0 & 0 \end{array}\right)\right) \\
&= \left(\widetilde{G}_1 \middle| \widetilde{G}_2\right) \left(Q_1 \middle| Q_2\right)^H \widehat{A} P_1 - \left(\widetilde{G}_1 \middle| \widetilde{G}_2\right) \left(Q_1 \middle| Q_2\right)^H \left(Q_1 \middle| Q_2\right) \left(\begin{array}{c|c} R_{11} & R_{12} \\ \hline 0 & 0 \end{array}\right) \quad (1)
\end{aligned}$$

We now *choose* the randomizing matrix $\left(\widetilde{G}_1 \middle| \widetilde{G}_2\right)$ for this step by setting

$$\left(\widetilde{G}_1 \middle| \widetilde{G}_2\right) = \left(G_1 \middle| G_2\right) \left(Q_1 \middle| Q_2\right). \quad (2)$$

Inserting the choice (2) into (1), we obtain

$$\left(\widetilde{Y}_1 \middle| \widetilde{Y}_2\right) = \underbrace{\left(G_1 \middle| G_2\right) \widehat{A}}_{\left(Y_1 \middle| Y_2\right)} P_1 - \left(G_1 \middle| G_2\right) \left(I - \left(\frac{U_{11}}{U_{21}}\right) T_{11}^{-1} \left(\frac{U_{11}}{U_{21}}\right)^H\right) \left(\begin{array}{c|c} R_{11} & R_{12} \\ \hline 0 & 0 \end{array}\right).$$

Letting $\left(\overline{Y}_1 \middle| \overline{Y}_2\right) = \left(Y_1 \middle| Y_2\right) P_1$ we conclude that

$$\begin{aligned}
\widetilde{Y}_2 &= \overline{Y}_2 - \left(G_1 \middle| G_2\right) \left(I - \left(\frac{U_{11}}{U_{21}}\right) T_{11}^{-1} \left(\frac{U_{11}}{U_{21}}\right)^H\right) \left(\frac{R_{12}}{0}\right) \\
&= \overline{Y}_2 - \left(G_1 \middle| G_2\right) \left(\left(\frac{R_{12}}{0}\right) - \left(\frac{U_{11}}{U_{21}}\right) T_{11}^{-1} {U_{11}}^H R_{12}\right) \\
&= \overline{Y}_2 - \left(G_1 - (G_1 U_{11} + G_2 U_{21}) T_{11}^{-1} {U_{11}}^H\right) R_{12}.
\end{aligned}$$

which can then be used to downdate the sampling matrix Y.

Remark 2 *Since the probability distribution for Gaussian random matrices is invariant under unitary maps, the formula (2) appears quite safe. After all,* $\left(G_1 \middle| G_2\right)$ *is Gaussian, and* $\left(Q_1 \middle| Q_2\right)$ *is just a sequence of reflections, so it might be tempting to conclude that the new randomizing matrix must be Gaussian too. However, the matrix* $\left(Q_1 \middle| Q_2\right)$ *unfortunately depends on the draw of the previous randomizing matrices, so this argument does not work. Nevertheless, the dependence of* $\left(Q_1 \middle| Q_2\right)$ *on the previous draws is very subtle since this matrix is dictated primarily by the directions of the good pivot columns. Extensive practical experiments (see, e.g., Section 4.2) indicate that the downdating strategy described in this section is just as good as the one that uses "pure" Gaussian matrices in Section 3.1.*

3.3 The blocked algorithm

In Fig. 4, we give the blocked algorithm that results when the randomized pivot selection strategy described in Section 3.1 is combined with the downdating techniques described in Section 3.2.

3.4 Asymptotic cost analysis

In analyzing the cost of the algorithm, we note that all steps in Fig. 4 highlighted in grey are part of the blocked HQR algorithm, which is known to have a cost of approximately $4/3n^3$ flops. This leaves us to discuss the overhead related to the other operations.

- $G := \text{RAND_IID}(b+p, n(A)))$: Cost: ignored.
- $Y := GA$: Cost: $2(b+p)n^2$ flops.
- $s_1 := \text{DeterminePivots}(\left(Y_1 \middle| Y_2 \right), b)$: Cost: $O(b(b+p)n)$ flops. (Recall that the factorization of this matrix can stop after the first b columns have been identified.)
- $\cdots := \text{SwapCols}(s_1, \cdots)$: Cost: ignored.
- $Y_2 := Y_2 - \left(G_1 - (G_1 U_{11} + G_2 U_{21}) T_{11}^{-1} {U_{11}}^H\right) R_{12}$: Aggregate cost: $O((b+p)n^2)$.

Thus, the overhead is $O((b+p)n^2)$ flops and the total cost is

$$4/3n^3 + O((b+p)n^2) \text{ flops},$$

which, asymptotically, approaches $4/3n^3$ flops. (A similar analysis can be performed for rectangular matrices.)

4 Experiments

This section describes the results from two sets of experiments. Section 4.1 compares the computational speed of the proposed scheme to existing state-of-the-art methods for computing HQR. Section 4.2 investigates how well the proposed randomized technique works at selecting pivot columns. Specifically, we investigate how well the rank-k truncated QR factorization approximates the original matrix, and compare the results to those obtained by classical column pivoting.

4.1 Performance experiments

We have implemented the proposed HQRRP algorithm using the `libflame` [31, 30] library that allows our implementations to closely resemble the algorithms as presented in the paper.

We would like to emphasize that the implementation we describe is not trying to push the limits of performance. It is merely trying to demonstrate that HQRRP is competitive and worthy of future study. More on this in the conclusion.

Platform details. All experiments reported in this article were performed on an Intel Xeon E5-2695 v3 (Haswell) processor (2.3 GHz), with 14 cores. In order to be able to show scalability results, the clock speed was throttled at 2.3 GHz, turning off so-called turbo boost. Each core can execute 16 double precision floating point operations per cycle, meaning that the peak performance of one core is 36.8 GFLOPS (billions of floating point operations per second). For comparison, on a single core, `dgemm` achieves around 33.6 GFLOPS. Other details of interest include that the OS used was Linux (Version 2.6.32-504.el6.x86_64), the code was compiled with gcc (Version 4.4.7), `dgeqrf` and `dgeqp3` were taken from LAPACK (Release 3.4.0), and the implementations were linked to BLAS from Intel's MKL library (Version 11.2.3). Our implementations were coded with `libflame` (Release 11104).

Implementations. We report performance for five implementations:

`dgeqrf`. The implementation of blocked HQR that is part of the `netlib` implemenation of LAPACK, modified so that the block size can be controlled.

`dgeqp3`. The implementation of blocked HQRP that is part of the `netlib` implemenation of LAPACK, based on [27], modified so that the block size can be controlled.

`HQRRP`. Our implementation of HQRRP that computes new matrices G and Y in every iteration. This implementation deviates from the algorithm in Fig. 4 in that it also incorporates additional column pivoting within the the call to HQRRP_UNB. For our experiments, we used an "oversampling" of $p = 10$.

`HQRRPddY`. The implementation of HQRRP that downdates Y. (It also includes pivoting within HQRRP_UNB).

`dgeqpx`. An implementation of HQRP with window pivoting, briefly mentioned in the introduction. This algorithm consists of two stages: The first stage is a QR with window pivoting. An incremental condition estimator is employed to accept/reject the columns within a window. The size of the window is about about twice the block size used by the algorithm to maintain locality. If some of the columns in the window are unacceptable, these are rejected and fresh ones are brought into the window. The second stage is a postprocessing stage to validate the rank, that is, to check that all good columns are in the front, and all the bad columns are in the rear. This step is required because the window pivoting could fail to reveal the rank due to its shortsightedness (having only checked the window and having employed a cheap condition estimator.) Sometimes, some columns must be moved between R_{11} that has been computed and matrices R_{12} and R_{22}, and then retriangularization must be performed with Givens rotations.

In all cases, we used an algorithmic block sizes of $b = 64$ and 128. While likely not optimal for all problem sizes, these blocks sizes yield near best performance and, regardless, it allows us to easily compare and contrast the performance of the different implementations.

Results. In Fig. 5, we report performance on one core and four cores. Performance is very encouraging. Not only does `HQRRPddY` handily outperform `dgeqp3` and to a lesser degree `dgeqpx`, but asymptotically its performance approaches that of `dgeqrf`. Parallel performance on four cores was attained via multithreaded BLAS, which is known to limit the opportunity for parallelism. Still, reasonable performance is attained. A complete analysis of how to parallelize these implementations and the resulting scalability is beyond the scope of this paper, which is why we do not report performance results for more cores.

4.2 Quality experiments

In this section, we investigate how good the pivot choices made by the proposed randomized technique are. We tested both the randomized scheme and classical column pivoted QR on four different test matrices:

- *Matrix 1 (fast decay):* This is an $n \times n$ matrix of the form $A = UDV^*$ where U and V are randomly drawn matrices with orthonormal columns (obtained by performing QR on a random Gaussian matrix), and where D is diagonal with entries given by $d_j = \beta^{(j-1)/(n-1)}$ with $\beta = 10^{-5}$.

- *Matrix 2 (S shaped decay):* This matrix is built in the same manner as "Matrix 1", but now the diagonal entries of D are chosen to first hover around 1, then decay rapidly, and then level out at 10^{-6}, as shown in Figure 7 (black line).

- *Matrix 3 (Gauss):* This matrix is a Gaussian random matrix scaled so that its leading singular value is 1. (To be precise, $A = \beta G$ where G is an $n \times n$ random matrix whose entries are drawn from an IID normalized Gaussian distribution, and $\beta = 1/\|G\|$.)

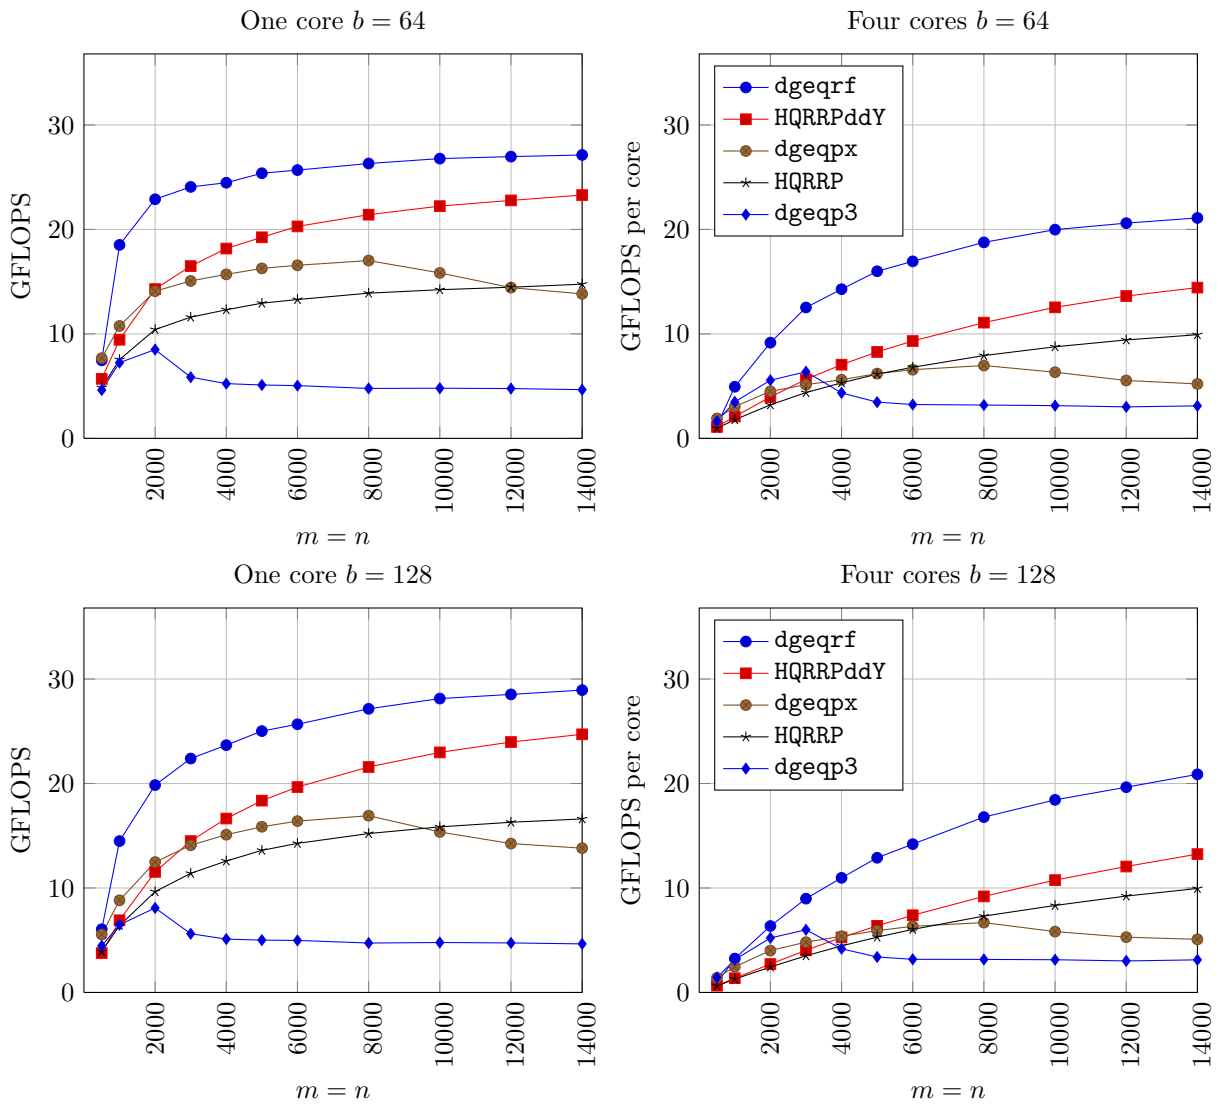

Figure 5: Performance of the various implementations. The top of the graphs represents theoretical peak (when "turbo boost" is disabled).

- *Matrix 4 (Kahan):* This is a variation of the "Kahan counter-example" [20] which is designed to trip up classical column pivoting. The matrix is formed as $A = SK$ where:

$$S = \begin{bmatrix} 1 & 0 & 0 & 0 & \cdots \\ 0 & \zeta & 0 & 0 & \cdots \\ 0 & 0 & \zeta^2 & 0 & \cdots \\ 0 & 0 & 0 & \zeta^3 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \quad \text{and} \quad K = \begin{bmatrix} 1 & -\phi & -\phi & -\phi & \cdots \\ 0 & 1 & -\phi & -\phi & \cdots \\ 0 & 0 & 1 & -\phi & \cdots \\ 0 & 0 & 0 & 1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

  for some positive parameters ζ and ϕ such that $\zeta^2 + \phi^2 = 1$. In our experiments, we chose $\zeta = 0.99$.

For each test matrix, we computed QR factorizations

$$AP = QR \tag{3}$$

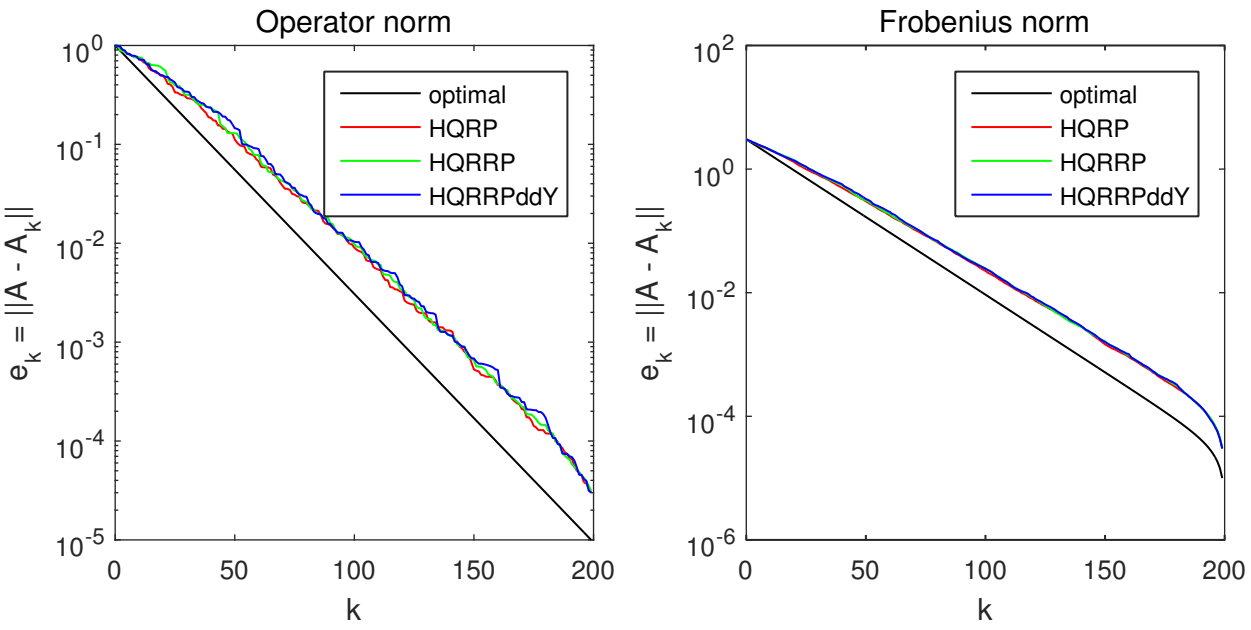

Figure 6: Residual errors e_k for "Matrix 1" as a function of the truncation rank k, cf. (4). The red line shows the results from traditional column pivoting, while the green and blue lines refer to the randomized methods we propose. The black line indicates the theoretically minimal errors resulting from a rank-k partial singular value decomposition.

using three different techniques:

- `HQRP`: The standard QR factorization `qr` built in to Matlab R2015a.
- `HQRRP`: The randomized algorithm described in Figure 4, but without the updating strategy for computing the sample matrix Y.
- `HQRRPddY`: The randomized algorithm described in Figure 4.

In all experiments, we used test matrices of size 200×200, a block size of $b = 20$, and an over-sampling parameter of $p = 5$.

As a quality measure of the pivoting strategy, we computed the errors e_k incurred when the factorization is truncated to its first k components. To be precise, these residual errors are defined via

$$e_k = \|AP - Q(:, 1:k)R(1:k, :)\| = \|R((k+1):n, (k+1):n)\|. \tag{4}$$

The results are shown in Figures 6 – 9, for the four different test matrices. The black lines in the graphs show the theoretically minimal errors incurred by a rank-k approximation. These are provided by the Eckart-Young theorem [9] which states that, with $\{\sigma_j\}_{j=1}^n$ denoting the singular values of A:

$$\begin{aligned} &e_k \geq \sigma_{k+1} && \text{when errors are measured in the spectral norm, and} \\ &e_k \geq \left(\textstyle\sum_{j=k+1}^n \sigma_j\right)^{1/2} && \text{when errors are measured in the Frobenius norm.} \end{aligned}$$

We observe in all cases that the quality of the pivots chosen by the randomized method very closely match those resulting from classical column pivoting (with the exception of the Kahan counter-example, where randomized pivot selection is better).

When classical column pivoting is used, the factorization (3) produced has the property that the diagonal entries of R are strictly decaying in magnitude

$$|R(1,1)| \geq |R(2,2)| \geq |R(3,3)| \geq \cdots$$

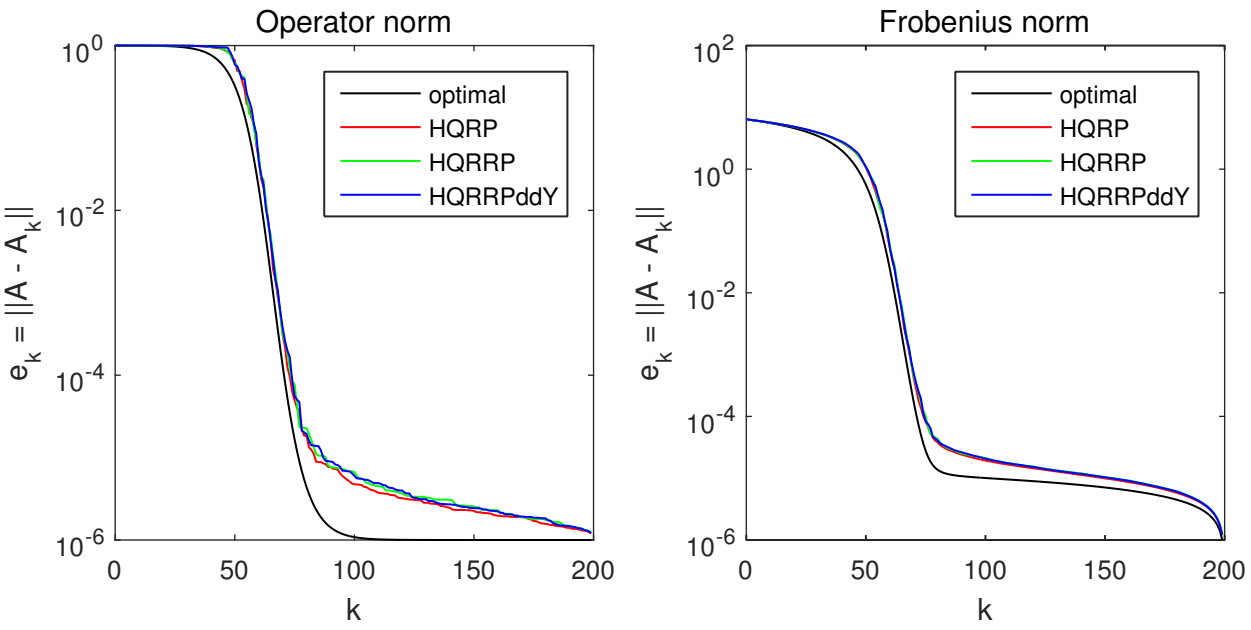

Figure 7: Residual errors e_k for "Matrix 2" as a function of the truncation rank k. Notation is the same as in Figure 6.

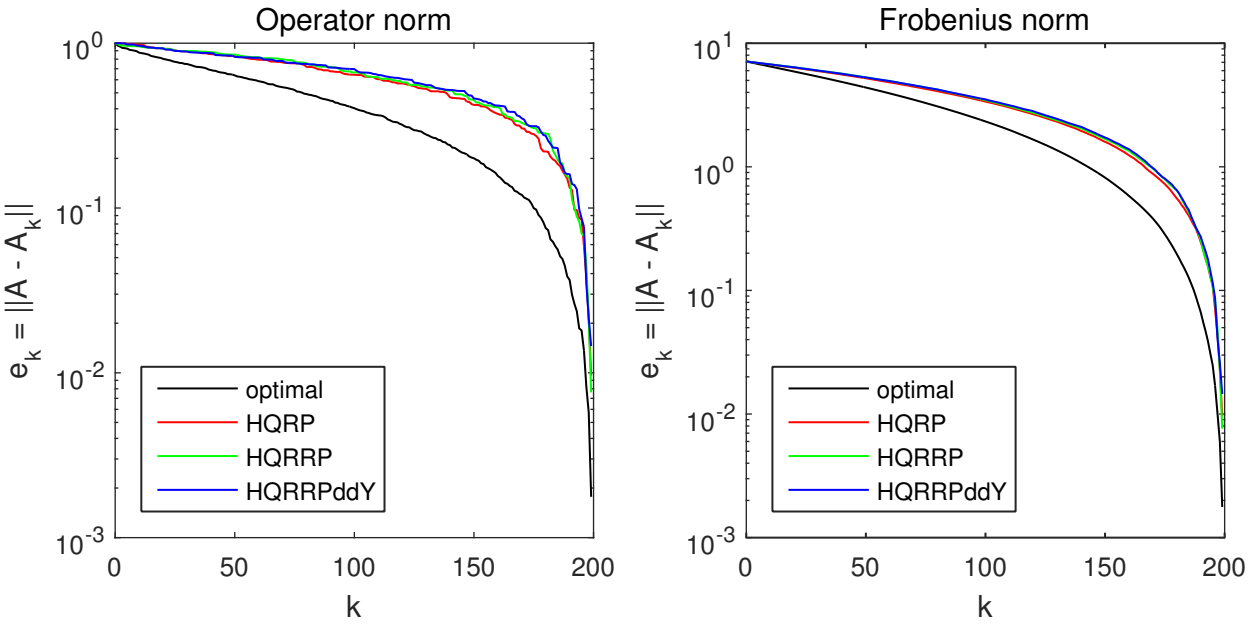

Figure 8: Residual errors e_k for "Matrix 3" as a function of the truncation rank k. Notation is the same as in Figure 6.

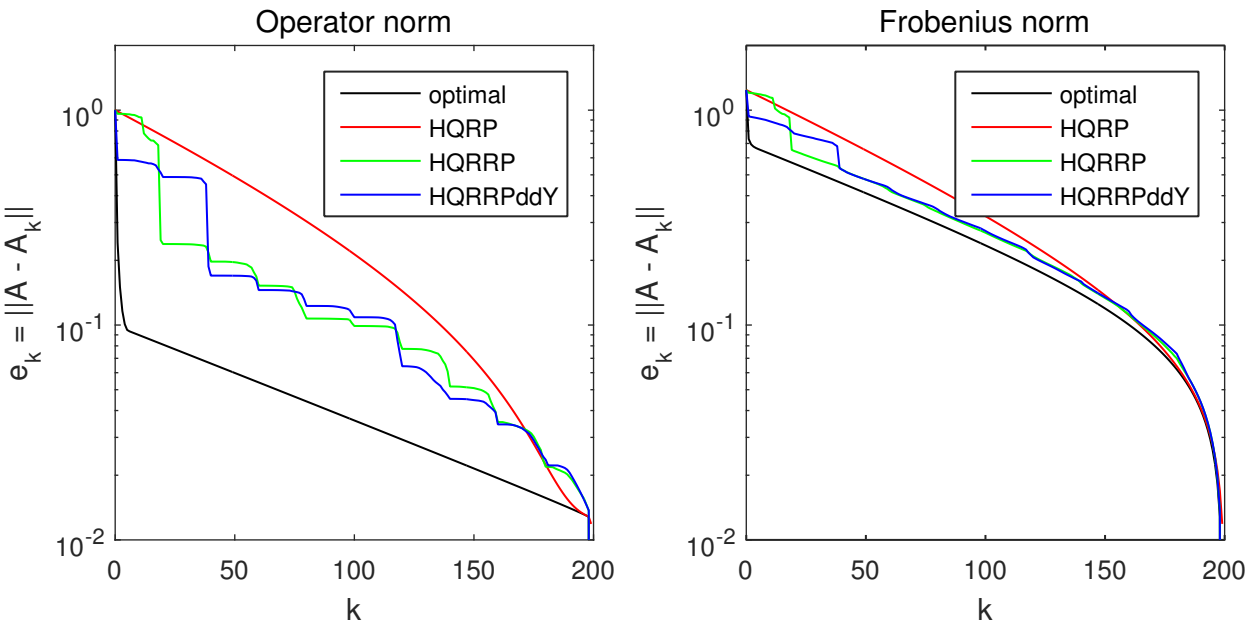

Figure 9: Residual errors e_k for "Matrix 4" as a function of the truncation rank k. Notation is the same as in Figure 6.

When the randomized pivoting strategies are used, this property is not enforced. To illustrate this point, we show in Figure 10 the values of the diagonal entries obtained by the randomized strategies versus what is obtained for classical column pivoting.

5 Conclusions and future work

We have described the algorithm HQRRP which is a blocked version of Householder QR with column pivoting. The main innovation compared to earlier work is that pivots are determined in groups by using randomized sampling. We demonstrated that the quality of the chosen pivots is comparable to traditional column pivoting, and that the dominant term in the asymptotic flop count equals that of non-pivoted QR. Importantly, we also demonstrate through numerical experiments that HQRRP is very fast, in fact almost as fast as unpivoted QR for a single-core implementation, and within a factor of 2 for a multi-core code. We believe that additional optimization should close the gap even further, in particular in the multi-core case.

The randomized technique described opens up several potential avenues for future research. The speed gains we demonstrate on single core and shared memory multicore machines are due to the reduction in data movement. Equivalently, data moved between memory layers is carefully amortized. We expect the technique described to have an even more pronounced advantage over traditional column pivoted QR when implemented in more severely communication constrained environments such as a matrix processed on a GPU or a distributed memory parallel machine, or stored out-of-core.

The randomized sampling techniques we describe can also be used to construct very close to optimal rank-revealing factorizations. To describe the idea, we note that a column pivoted QR factorization of a given matrix A can be written as

$$A = Q\,R\,P^*, \tag{5}$$

where Q is orthonormal, R is upper triangular, and P is a permutation matrix. In this manuscript, we used randomized sampling to determine the permutation matrix P. It turns out that for a modest additional cost, one can build a factorization that takes the form (5), but with both Q and P built as products of Householder reflectors. This generalization allows us to bring R not only to upper triangular form, but

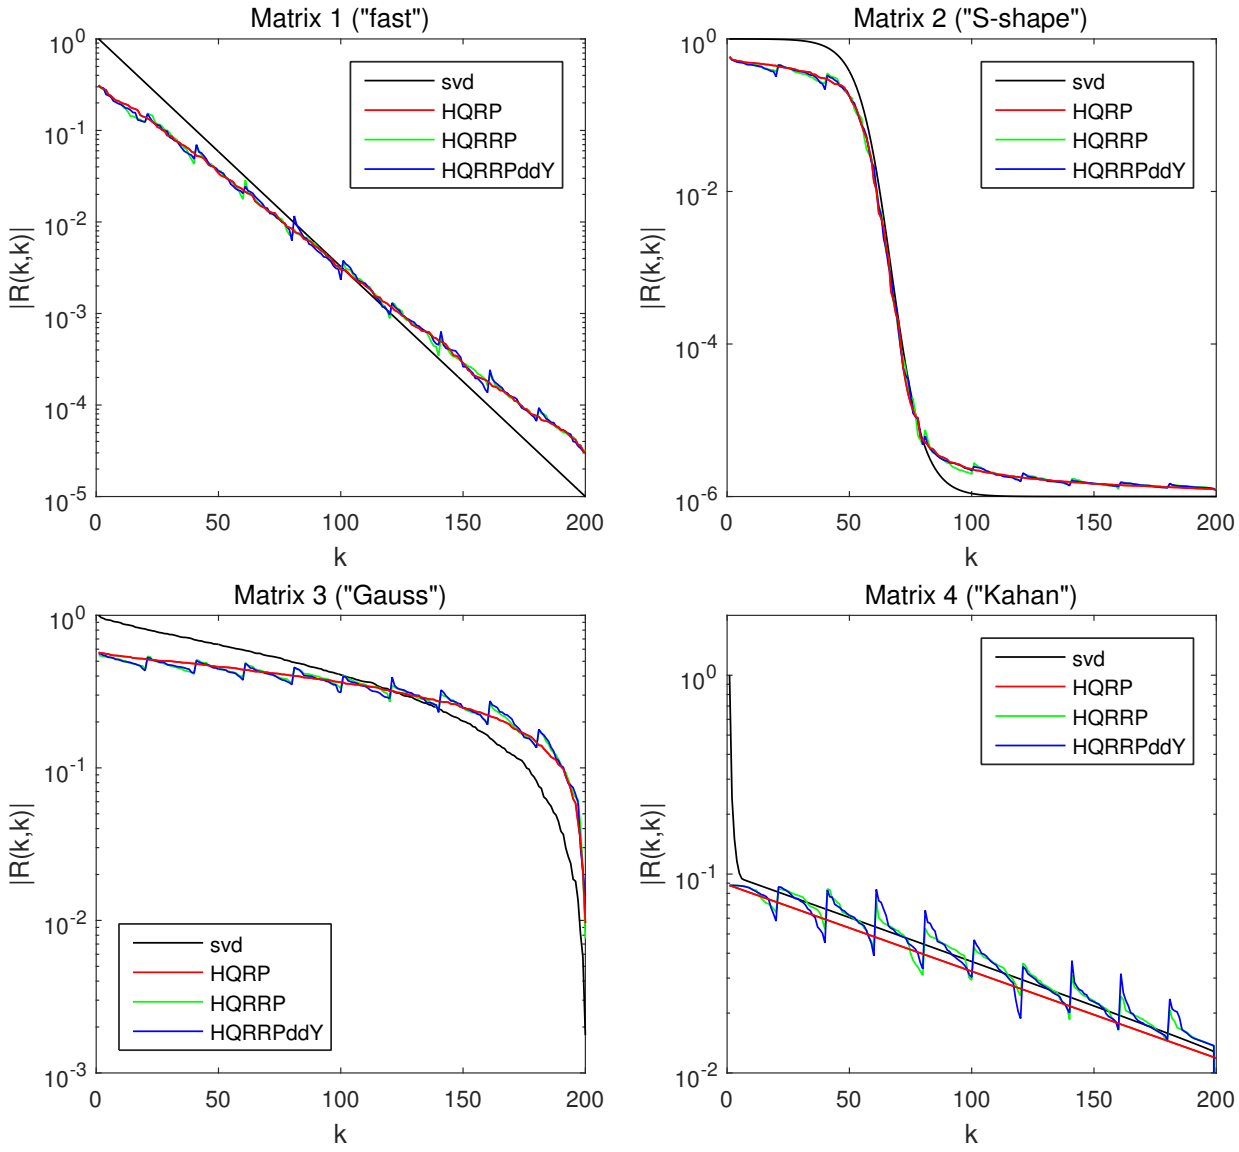

Figure 10: For each of the four test matrices described in Section 4.2, we show the magnitudes of the diagonal entries in the "R"-factor in a column pivoted QR factorization. We compare classical column pivoting (red) with the two randomized techniques proposed here (blue and green). We also show the singular values of each matrix (black) for reference.

very close to being diagonal, with accurate approximations to the singular values of A on its diagonal. This discovery was reported in [22], and is a subject of ongoing research.

How to scale the presented algorithm to very large numbers of cores is an open research question. Techniques such as "compute ahead" will have to be employed to ensure that the factorization of the current panel (A_{11} and A_{21}) and downdate of Y do not start dominating the parts of the computation that can be cast in terms of GEMM.

Acknowledgments

The research reported was supported by DARPA, under the contract N66001-13-1-4050, and by the NSF, under the contracts DMS-1407340 and ACI-1148125/1340293. We thank Francisco Igual for providing access to the equipment used for the performance experiments.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).

References

[1] AMD. AMD Core Math Library. `http://developer.amd.com/tools-and-sdks/cpu-development/amd-core-math-library-acml/`.

[2] E. Anderson, Z. Bai, C. Bischof, L. S. Blackford, J. Demmel, Jack J. Dongarra, J. Du Croz, S. Hammarling, A. Greenbaum, A. McKenney, and D. Sorensen. *LAPACK Users' guide (third ed.)*. SIAM, Philadelphia, PA, USA, 1999.

[3] C. Bischof and C. Van Loan. The WY representation for products of Householder matrices. *SIAM J. Sci. Stat. Comput.*, 8(1):s2–s13, Jan. 1987.

[4] C.H. Bischof and G. Quintana-Ortí. Algorithm 782: Codes for rank-revealing QR factorizations of dense matrices. *ACM Trans. Math. Soft.*, 24(2):254–257, 1998.

[5] C.H. Bischof and G. Quintana-Ortí. Computing rank-revealing QR factorizations of dense matrices. *ACM Trans. Math. Soft.*, 24(2):226–253, 1998.

[6] J.J. Dongarra, J. Du Croz, S. Hammarling, and I. Duff. A set of level 3 basic linear algebra subprograms. *ACM Trans. Math. Soft.*, 16(1):1–17, March 1990.

[7] J.J. Dongarra, I.S. Duff, D.C. Sorensen, and H. A. van der Vorst. *Solving Linear Systems on Vector and Shared Memory Computers*. SIAM, Philadelphia, PA, 1991.

[8] J. Duersch and M. Gu. True BLAS-3 performance QRCP using random sampling, 2015. arXiv preprint #1509.06820.

[9] C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.

[10] A. Frieze, R. Kannan, and S. Vempala. Fast Monte-Carlo algorithms for finding low-rank approximations. *J. ACM*, 51(6):1025–1041 (electronic), 2004.

[11] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, 3nd edition, 1996.

[12] K. Goto and R.A. van de Geijn. High-performance implementation of the level-3 BLAS. *ACM Trans. Math. Softw.*, 35(1):4:1–4:14, July 2008.

[13] Kazushige Goto and Robert van de Geijn. Anatomy of high-performance matrix multiplication. *ACM Trans. Math. Soft.*, 34(3):12:1–12:25, May 2008.
Download from http://www.cs.utexas.edu/users/flame/web/FLAMEPublications.html.

[14] M. Gu and S. C. Eisenstat. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM J. Sci. Comput.*, 17(4):848–869, 1996.

[15] J.A. Gunnels, F.G. Gustavson, G.M Henry, and R.A. van de Geijn. FLAME: Formal Linear Algebra Methods Environment. *ACM Trans. Math. Soft.*, 27(4):422–455, December 2001.
Download from http://www.cs.utexas.edu/users/flame/web/FLAMEPublications.html.

[16] N. Halko, P.G. Martinsson, and J.A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.

[17] IBM. Engineering and Scientific Subroutine Library. `http://www-03.ibm.com/systems/power/software/essl/`.

[18] Intel. Math Kernel Library. `http://developer.intel.com/software/products/mkl/`.

[19] T. Joffrain, T. M. Low, E. S. Quintana-Ortí, R. van de Geijn, and F. G. Van Zee. Accumulating Householder transformations, revisited. *ACM Trans. Math. Softw.*, 32(2):169–179, June 2006.

[20] W. Kahan. Numerical linear algebra. *Canadian Math. Bull*, 9(6):757–801, 1966.

[21] E. Liberty, F. Woolfe, P.G. Martinsson, V. Rokhlin, and M. Tygert. Randomized algorithms for the low-rank approximation of matrices. *Proc. Natl. Acad. Sci. USA*, 104(51):20167–20172, 2007.

[22] P.G. Martinsson. Blocked rank-revealing QR factorizations: How randomized sampling can be used to avoid single-vector pivoting, 2015. arXiv preprint #1505.08115.

[23] P.G. Martinsson, V. Rokhlin, and M. Tygert. A randomized algorithm for the approximation of matrices. Technical Report Yale CS research report YALEU/DCS/RR-1361, Yale University, Computer Science Department, 2006.

[24] P.G. Martinsson, V. Rokhlin, and M. Tygert. A randomized algorithm for the decomposition of matrices. *Appl. Comput. Harmon. Anal.*, 30(1):47–68, 2011.

[25] C. Puglisi. Modification of the Householder method based on the compact WY representation. *SIAM J. Sci. Stat. Comput.*, 13:723–726, 1992.

[26] E.S. Quintana, G. Quintana, X. Sun, and R. van de Geijn. A note on parallel matrix inversion. *SIAM J. Sci. Comput.*, 22(5):1762–1771, 2001.

[27] G. Quintana-Ort, X. Sun, and C.H. Bischof. A BLAS-3 version of the QR factorization with column pivoting. *SIAM Journal on Scientific Computing*, 19(5):1486–1494, 1998.

[28] R. Schreiber and C. Van Loan. A storage-efficient WY representation for products of Householder transformations. *SIAM J. Sci. Stat. Comput.*, 10(1):53–57, Jan. 1989.

[29] G.W. Stewart. *Matrix Algorithms Volume 1: Basic Decompositions*. SIAM, 1998.

[30] F.G. Van Zee. `libflame`*: The Complete Reference*. `www.lulu.com`, 2012.
Download from http://www.cs.utexas.edu/users/flame/web/FLAMEPublications.html.

[31] F.G. Van Zee, E. Chan, R. van de Geijn, E.S. Quintana-Ortí, and G. Quintana-Ortí. The libflame library for dense matrix computations. *IEEE Computation in Science & Engineering*, 11(6):56–62, 2009.

[32] F.G. Van Zee and R.A. van de Geijn. BLIS: A framework for rapidly instantiating BLAS functionality. *ACM Transactions on Mathematical Software*, 41(3), 2015.

[33] S. Voronin and P.G. Martinsson. A CUR factorization algorithm based on the interpolative decomposition. *arXiv.org*, 1412.8447, 2014.

[34] H. F. Walker. Implementation of the GMRES method using Householder transformations. *SIAM J. Sci. Stat. Comput.*, 9(1):152–163, 1988.

[35] X. Zhang, Q. Wang, and Y. Zhang. Model-driven level 3 BLAS performance optimization on Loongson 3A processor. In *Proceedings of the 2012 IEEE 18th International Conference on Parallel and Distributed Systems*, pages 684–691. IEEE Computer Society, 2012.